CS 880: Quantum Information Processing 11/1/10

Lecture 23: Cryptographic Protocols

Instructor: Dieter van Melkebeek Scribe: Cong Han Lim

In today and tomorrow's lectures we will be looking at cryptographic protocols. The protocols enable Alice and Bob to communicate with each other in a way where we can ensure the following goals.

1. Authenticity: Alice can authenticate a message she sends such that when Bob receives it, Bob knows that it is from Alice.
2. Secrecy: An eavesdropper (Eve) on the communication channel cannot decode the secret message.

To achieve these goals we will be looking at some basic primitives. In this lecture we will cover secret key exchange and bit commitment, and tomorrow we will be talking about zero knowledge protocols.

We had already encountered the concept of secret key exchange when we talked about the Diffie-Hellman key exchange. The goal there was for Alice and Bob to share a randomly chosen secret key, which can subsequently be used for encryption and decryption. For security reasons, only Alice and Bob should know the key, and the domain the key was chosen from has to be sufficiently large to prevent simple exhaustive search attacks.

As for bit commitment, the goal is to allow Alice to commit to a value without revealing it. The real world analogue would be Alice sending a message inside a locked box to Bob while she retains the key to it.

In the classical setting, we do not know how to achieve unconditonal security when implementing these primitives. Instead, we only know how to implement all of these primitives under certain hardness assumptions like relying on one-way functions. In principle, this means that the message being transmitted actually contains all the information necessary to decode it, though it is computationally difficult to break with a classical computer.

Things change once we are in the quantum setting. On the one hand, the list of problems that are "computationally difficult" changes once we have access to a quantum computer. For example, the Diffie-Hellman protocol, El-Gamal cryptosystems and RSA cryptosystems can be broken by quantum algorithms. On the other hand, a quantum computer can actually aid cryptography. As we shall see in this lecture, it allows us to realize secret key exchange unconditionally in an information-theoretic sense. However, we will also prove that it is impossible to unconditionally realize bit commitment in the same sense.

1 Secret Key Exchange

We will describe two protocols. The first one will be covered in detail. We will also briefly describe another (BB84) which is easier to implement and has in fact been realized physically (albeit with easy possible side-channel attacks). Both protocols capitalize on entanglement between the halves of EPR pairs, and presume a classical public channel that cannot be tampered with.

1.1 First Protocol (by Lo & Chau)

This protocol utilizes the key property that for an EPR pair where Alice and Bob have one part each, when Alice and Bob measure their respective qubits, they will both observe the same random bit. This is a way to share a single bit of a secret key, but it presupposes that there is some precommunication: the EPR pair has to generated by one party, split up and one qubit has to be sent to the other party, and during the transmission one could tamper with the qubit. We will describe how to get around this issue. We will be using the Bell basis throughout this protocol:

$$\begin{aligned}
\left|\Phi^+\right\rangle &= \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle),\\
\left|\Phi^-\right\rangle &= \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle),\\
\left|\Psi^+\right\rangle &= \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle),\\
\left|\Psi^-\right\rangle &= \frac{1}{\sqrt{2}}(|01\rangle - |10\rangle).
\end{aligned}$$

There are three steps in this protocol. First, Alice will generate $|\Phi^+\rangle^{\otimes n}$ (the tensor product of n EPR pairs) and send Bob his half of the state. Then both Alice and Bob will check the qubits they have to determine if the state is equal or close to $|\Phi^+\rangle^{\otimes n}$ or if it has been tampered with. If the states are perfect, the check should always accept. Otherwise, if the state is not close to $|\Phi^+\rangle^{\otimes n}$, then we want the check to return negative, after which Alice and Bob will discard the qubits and restart the protocol. After a successful check, Alice and Bob will measure the remaining components to obtain the secret key, which will be (almost) uniformly distributed.

How are we going to perform the checking process? Alice and Bob could measure some of the qubits directly, which would actually destroy them but would leave enough of the joint state to be used in obtaining the random key. To see how direct measurement helps, we first consider a single two qubit state in the Bell basis. To detect whether there is a high weight on the $|\Psi\rangle$ terms, Alice and Bob both measure their qubit and compare their measurement outcomes using the secure public channel. They reject if the outcomes differ, and continue the protocol otherwise. Conceptually, this step can be thought of as performing a $CNOT$ on the pair and observing the first qubit (the non-control qubit).

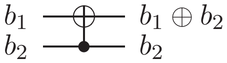

In the case of the $|\Phi\rangle$ states, the first qubit becomes $|0\rangle$, whereas for the $|\Psi\rangle$ states we get $|1\rangle$. Thus, if we measure that qubit and observe a 1, we can reject. Otherwise, we know that our state now only has $|\Phi\rangle$ components.

The above procedure allows us to eliminate the unwanted $|\Psi\rangle$ components. What about $|\Phi^-\rangle$? Applying the Hadamard gate to both qubits of our base states helpfully preserves $|\Phi^+\rangle$ but switches $|\Phi^-\rangle$ and $|\Psi^+\rangle$:

$$\begin{aligned}
(H \otimes H)\left(\left|\Phi^+\right\rangle\right) &= \left|\Phi^+\right\rangle,\\
(H \otimes H)\left(\left|\Phi^-\right\rangle\right) &= \left|\Psi^+\right\rangle,\\
(H \otimes H)\left(\left|\Psi^+\right\rangle\right) &= \left|\Phi^-\right\rangle,\\
(H \otimes H)\left(\left|\Psi^-\right\rangle\right) &= \left|\Psi^-\right\rangle.
\end{aligned}$$

We can now apply the $CNOT$ and observe the first qubit to eliminate the $|\Phi^-\rangle$ state. So, depending on whether we apply the Hadamard gate before the $CNOT$, we have two separate tests that allow us to eliminate the unwanted states. Alice picks one of the two tests uniformly at random and sends a bit over the public channel to Bob to coordinate the check. For this combined test we have that

$$\begin{aligned}\Pr[\left|\Phi^+\right\rangle \text{ passes}] &= 1\\ \Pr[|\Psi\rangle \text{ passes}] &\leq \left|\left\langle\Phi^+|\Psi\right\rangle\right|^2 + \frac{1}{2}\left(1 - \left|\left\langle\Phi^+|\Psi\right\rangle\right|^2\right)\\ &= \frac{1}{2}\left(1 + \left|\left\langle\Phi^+|\Psi\right\rangle\right|^2\right).\end{aligned}$$

The first equality follows because $|\Phi^+\rangle$ passes both tests. The inequality follows because the Bell states other than $|\Phi^+\rangle$ are rejected by at least one of the two tests.

We now consider the entire $2n$ qubit state. This can be seen as a superposition over all combinations of n tensors of the Bell basis. The direct approach would be to perform the check on some of the pairs of qubits, but the probability of detecting tampering if Eve only tampers with a small number of pairs is low. We need to use a better approach to increase the probability.

View the situation as follows. There are n possible combined tests Alice and Bob can perform, namely one for each of their n paired qubits. Each of those tests involves comparing two classical bits, one from Alice and one from Bob, which are obtained by measuring the corresponding qubits. Each individual test passes iff the bits are the same.

We'd like to design a new test that rejects with high probability if at least one of those individual tests rejects, and accepts otherwise. We can do so by taking a random subset of the individual tests and reject iff the parity of the number of those that reject is odd. If all individual tests accept, the parity is always even; otherwise, it is odd with probability 50%. Moreover, Alice and Bob only need to sacrifice one qubit pair in order to execute this new test. They use the public channel to select the subset, and then XOR into one of the selected positions (say the first one) all the other selected qubits. They then measure the first selected position, compare the measurement outcomes using the public channel, and reject if they differ.

Let us analyze what happens to the state of the system when we apply several new tests. We can write the initial state as a superposition over all possible tensors of Bell states. Consider the evolution without renormalization. The weight of the component $|\Phi^+\rangle^{\otimes n}$ remains unaffected as that state passes all tests. The expected weight of all other components decreases by a factor of $1/2$ with each new test. If the initial weight of $|\Phi^+\rangle^{\otimes n}$ is very small, then chances are at least one of the new tests rejects. Otherwise, the tests will make the expected weight of the other components small compared to the weight of the $|\Phi^+\rangle^{\otimes n}$ component, so in the end the normalized state is very close to $|\Phi^+\rangle^{\otimes n}$. Somewhat more quantitatively, if we haven't rejected after $n/2$ new tests, the remaining $n/2$ qubits are exponentially close to $|\Phi^+\rangle^{\otimes n}$ with exponentially high confidence.

One subtle concern we have yet to address is whether the information broadcast over the public channel can allow Eve to obtain any additional knowledge about the key. Since the only information we send over th subset, the type of check we are performing, and the results of the measurements, one can see that Eve cannot benefit from this additional knowledge.

1.2 Second Protocol (BB84)

While the first protocol gives us a fairly efficient way of performing information-theoretically secure quantum key exchange, it is not that easy to implement. In particular, implementing the XOR process is complicated as it is a quantum operation that involves many qubits. In comparison, the design and implementation of the very first quantum key exchange protocol (BB84) introduced by Charles Bennett and Gilles Brassard is rather elegant, though it took about fifteen years before the protocol was proven to be information-theoretically secure. We will briefly describe the BB84 protocol here.

For the initial communication over a quantum channel:

1. Alice randomly generates two binary strings $a = a_1 a_2 \ldots a_n$ and $b = b_1 b_2 \ldots b_n$ of length n. String a will form the basis of the secret key, and string b will be used to decide how we encode a as qubits. If b_i is 0 we encode a_i in the standard $\{|0\rangle, |1\rangle\}$ basis, otherwise we encode a_i in the Hadamard basis $\{|+\rangle, |-\rangle\}$.

2. Alice sends the encoded qubits to Bob, who has no way of knowing which basis was used for each qubit. Instead, Bob generates the random string b' and use the i-th bit of b' to decide which basis he is going to measure the i-th qubit in. This gives Bob another string a'. We know for sure that the i-th bit of a' is the same as a_i if b' and b are the same on the corresponding bit.

Since the choice of b' is random, Bob is expected to get at least half of b correct. Now we want to decide on our secret key. This is done over the public channel:

1. Alice and Bob share their b and b' over the public channel.

2. The bits of a and a' where b and b' do not coincide are discarded.

If Eve has not observed or tampered with the encoded qubits, the resulting a and a' will be the same; otherwise they may not be in total agreement. Alice and Bob now check if tampering has occured:

1. Alice and Bob disclose a portion of their strings a and a' over the public channel. These bits are removed from the secret key.

2. If there are any differences, then we know that Eve has been tampering or observing the state, and the protocol is aborted.

3. Otherwise, the remaining bits in a and a' form (two copies of) the secret key.

While this protocol is simpler to implement, the correctness proof is more difficult. One needs to argue that the more information an eavesdropper has extracted from the channel, the more likely it is for more components of a and a' to differ.

One disadvantage shared by both protocols is that they do not work in an on-line fashion, since one needs to send over the entire key before any processing can occur. There are other protocols that address this issue.

2 Bit Commitment

Bit commitment is a two-stage protocol for commiting to a value without revealing it. In the commitment phase, Alice picks a value for a bit b and sends some information to Bob that commits her to this value of b without revealing it. The next phase is the reveal phase, where Alice sends some additional information to Bob that allows him to figure out the bit Alice committed to. We want this protocol to satisfy two properties:

1. Binding: Once Alice has committed to a bit, she cannot change it afterwards.
2. Hiding: The information that Bob receives in the first phase cannot be used to obtain the bit before the reveal phase.

In the classical setting, we can again realize bit commitment in a hardness sense by using one-way functions. One would hope that the quantum setting allows us to achieve this in an information-theoretical setting, but that unfortunately is not true.

Theorem 1. *There is no information-theoretically secure quantum bit commitment scheme that has both the binding and hiding properties.*

Proof. The proof is an application of the Schmidt decomposition.

Consider the state after the commitment phase. The state can be described by a density operator. Consider purifications $\left|\psi_{AB}^{(b)}\right\rangle$ for $b \in \{0,1\}$, where b is the bit that is committed by Alice.

Note that information-theoretic hiding means that Bob's reduced density operator is exactly the same in both cases:

$$\varrho_B^{(0)} = \varrho_B^{(1)}. \tag{1}$$

By appyling the Schmidt decomposition, we get orthonormal bases $\left\{\left|\phi_{A_i}^{(b)}\right\rangle\right\}$ for Alice's part of the state and $\left\{\left|\phi_{B_i}^{(b)}\right\rangle\right\}$ for Bob's part of the state, such that we can write

$$\left|\psi_{AB}^{(b)}\right\rangle = \sum_i \lambda_i^{(b)} \left|\phi_{A_i}^{(b)}\right\rangle \left|\phi_{B_i}^{(b)}\right\rangle,$$

and

$$\varrho_B^{(b)} = \sum_i \left(\lambda_i^{(b)}\right)^2 \left|\phi_{B_i}^{(b)}\right\rangle \left\langle \phi_{B_i}^{(b)}\right|.$$

Becauase of (1) we can assume that $\lambda_i^{(0)} = \lambda_i^{(1)} \doteq \lambda_i$ and $\left|\phi_{B_i}^{(0)}\right\rangle = \left|\phi_{B_i}^{(1)}\right\rangle \doteq |\phi_{B_i}\rangle$, so we can rewrite the purifications as

$$\left|\psi_{AB}^{(b)}\right\rangle = \sum_i \lambda_i \left|\phi_{A_i}^{(b)}\right\rangle |\phi_{B_i}\rangle.$$

Since both $\left|\phi_{A_i}^{(b)}\right\rangle$ are orthonormal bases over the same space, Alice can simply apply some local unitary transformation that transforms $\left|\phi_{A_i}^{(0)}\right\rangle$ to $\left|\phi_{A_i}^{(1)}\right\rangle$. This allows Alice to cheat and modify the bit, so our protocol fails to be binding. □

3 References

The first quantum key exchange method is described in Lo and Chau's article [2], which contains a good amount of exposition. The BB84 protocol [1] is the first quantum key exchange protocol designed and the first proof of its security can be found in [3].

References

[1] C. Bennett and G. Brassard Quantum Cryptography: Public key distribution and coin tossing. *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing (Bangalore, India, Dec.)*, pages 175 - 179, 1984.

[2] H. Lo and H. Chau. Unconditional Security of Quantum Key Distribution over Arbitrarily Long Distances. *arXiv:quant-ph/9803006v5*, 1999.

[3] D. Mayers. Unconditional security in quantum cryptography. *Journal of the ACM, vol. 48, no. 3*, Pages 351 - 406. 2001.